

# Open Transactions on Shared Memory

Marino Miculan

`marino.miculan@uniud.it`

Marco Peressotti

`marco.peressotti@uniud.it`

Andrea Toneguzzo

`toneguzzo.andrea@spes.uniud.it`

Laboratory of Models and Applications of Distributed Systems  
Department of Mathematics and Computer Science  
University of Udine, Italy

## Abstract

*Transactional memory* has arisen as a good way for solving many of the issues of lock-based programming. However, most implementations admit *isolated* transactions only, which are not adequate when we have to coordinate *communicating* processes. To this end, in this paper we present *OCTM*, an Haskell-like language with *open* transactions over shared transactional memory: processes can *join* transactions at runtime just by accessing to shared variables. Thus a transaction can co-operate with the environment through shared variables, but if it is rolled-back, also all its effects on the environment are retracted. For proving the expressive power of *OCTM* we give an implementation of *TCCS<sup>m</sup>*, a CCS-like calculus with open transactions.

## 1 Introduction

Coordination of concurrent programs is notoriously difficult. Traditional fine-grained lock-based mechanisms are deadlock-prone, inefficient, not composable and not scalable. For these reasons, *Software Transactional Memory* (STM) has been proposed as a more effective abstraction for concurrent programming [1, 9, 18]. The idea is to mark blocks of code as “atomic”; at runtime, these blocks are executed so that the well-known ACID properties are guaranteed. Transactions ensure deadlock freedom, no priority inversion, automatic roll-back on exceptions or timeouts, and greater parallelizability. Among other implementations, we mention *STM Haskell* [7], which allows atomic blocks to be composed into larger ones. STM Haskell adopts an *optimistic* evaluation strategy: the blocks are allowed to run concurrently, and eventually if an interference is detected a transaction is *aborted* and its effects on the memory are rolled back.

However, standard ACID transactions are still inadequate when we have to deal with *communicating* processes, i.e., which can exchange information *during* the transactions. This is very common in concurrent distributed programming, like in service-oriented architectures, where processes dynamically combine to form a transaction, and all have to either commit or abort together. In this scenario the participants cannot be enclosed in one transaction beforehand, because transactions are formed at runtime. To circumvent this issue, various forms of *open transactions* have been proposed, where the Isolation requirement is relaxed [2–4, 11, 13]. In particular, *TransCCS* and *TCCS<sup>m</sup>* are two CCS-like calculi recently introduced to model communicating transactions [4, 5, 11]. These calculi offer methodologies for proving important properties, such as fair-testing for proving liveness and bisimulations for proving contextual equivalences.

Now, if we try to implement cross-transaction communications *a la* *TCCS<sup>m</sup>* in STM Haskell or similar languages, it turns out that isolated transactions are not expressive enough. As an example, let us consider two *TCCS<sup>m</sup>* transactions  $\langle \bar{c}.P \blacktriangleright 0 \rangle \mid \langle c.Q \blacktriangleright 0 \rangle$  synchronizing on a

channel  $c$ . Following the standard practice, we could implement this synchronization as two parallel processes using a pair of semaphores  $c1, c2$  (which are easily realized in STM Haskell):

$\langle \bar{c}.P \triangleright 0 \rangle = \text{atomic } \{$ up $c1$ -- 1.1 down $c2$ -- 1.2 P }	$\langle c.Q \triangleright 0 \rangle = \text{atomic } \{$ down $c1$ -- 2.1 up $c2$ -- 2.2 Q }
--	--

This implementation is going to deadlock: the only possible execution order is 1.1-2.1-2.2-1.2, which is possible outside transactions but it is forbidden for ACID transactions<sup>1</sup>. The problem is that ordinary STM transactions are kept isolated, while in  $TCCS^m$  they can merge at runtime.

In order to address this issue, in this paper we introduce software transactional memory with *open* transactions: processes can *join* transactions and transactions can *merge* at runtime, when they access to shared variables. To this end, we present *OCTM*, a higher-order language extending the concurrency model of STM Haskell with composable *open (multi-thread)* transactions interacting via *shared memory*. The key step is to separate the isolation aspect from atomicity: in *OCTM* the `atomic` construct ensures “all-or-nothing” execution, but not isolation; when needed, isolated execution can be guaranteed by a new constructor `isolated`. An `atomic` block is a *participant* (possibly the only one) of a transaction. Notice that transaction merging is implicitly triggered by accessing to shared memory, without any explicit operation or *a priori* coordination. For instance, in *OCTM* the two transactions of the example above would merge becoming two participants of the same transaction, hence the two threads can synchronize and proceed. In order to prove formally the expressivity of open memory transactions, we define an implementation of  $TCCS^m$  in *OCTM*, which is proved to correctly preserve behaviours by means of a suitable notion of simulation. We have based our work on STM Haskell as a paradigmatic example, but this approach is general and can be applied to other STM implementations.

Lesani and Palsberg [13] have proposed transactions communicating through transactional message-based channels called *transactional events*. These mechanisms are closer to models like TransCCS and  $TCCS^m$ , but on the other hand they induce a *strict coupling* between processes, which sometimes is neither advisable nor easy to implement (e.g., when we do not know all transaction’s participants beforehand). In fact, most STM implementations (including STM Haskell) adopt the shared memory model of multi-thread programming; this model is also more amenable to implementation on modern multi-core hardware architectures with transactional memory [8]. For these reasons, in *OCTM* we have preferred to stick to *loosely coupled* interactions based on shared memory only.

The rest of the paper is structured as follows. In Section 2 we describe the syntax and semantics of *OCTM*. Some examples are in Section 3. In Section 4 we assess the expressiveness of *OCTM* by providing an implementation of  $TCCS^m$ , our reference model for open transactions. Conclusions and directions for future work are in Section 5. Longer proofs are in the Appendix.

## 2 OCTM: Open Concurrent Transactional Memory

In this section we introduce the syntax and semantics of *OCTM*, a higher-order functional language with threads and open transaction on shared memory. The syntax is Haskell-like (in the wake of existing works on software transactional memories such as [7]) and the semantics is a small-step operational semantics given by two relations:  $\xrightarrow{\beta}$  models transaction auxiliary opera-

---

<sup>1</sup>This possibility was pointed out also in [7]: “two threads can easily deadlock if each awaits some communication from the other”.

Value	$V ::= r \mid \lambda x. M \mid \text{return } M \mid M \gg= N \mid$ $\text{newVar } M \mid \text{readVar } r \mid \text{writeVar } r \ M \mid$ $\text{fork } M \mid \text{atomic } M \ N \mid \text{isolated } M \mid \text{abort } M \mid \text{retry}$
Term	$M, N ::= x \mid V \mid MN \mid \dots$

Figure 1: Syntax of *OCTM* values and terms.

tions (e.g. creation) while  $\rightarrow$  models actual term evaluations. Executions proceeds by repeatedly choosing a thread and executing a single (optionally transactional) operation; transitions from different threads may be arbitrarily interleaved as long as atomicity and isolation are not violated where imposed by the program.

## 2.1 Syntax

The syntax can be found in Figure 1 where the meta-variables  $r$  and  $x$  range over a given countable set of locations  $\text{Loc}$  and variables  $\text{Var}$  respectively. Terms and values are inspired to Haskell and are entirely conventional<sup>2</sup>; they include abstractions, application, monadic operators (`return` and `>>=`), memory operators (`newVar`, `readVar`, `writeVar`), forks, transactional execution modalities (`atomic` and `isolated`) and transaction operators (`abort` and `retry`).

Effectfull expressions such as `fork` or `isolated` are glued together by the (overloaded) monadic bind `>>=` e.g.:

```
newVar 0 >>=  $\lambda x. (\text{fork} (\text{writeVar } x \ 42) >>= \lambda y. \text{readVar } x)$ 
```

whereas values are “passed on” by the monadic unit `return`.

Akin to Haskell, we will use underscores in place of unused variables (e.g.  $\lambda \_. 0$ ) and  $M \gg N$  as a shorthand for  $M \gg= \lambda \_. N$ , and the convenient *do-notation*:

```
do{ $x \leftarrow M; N$ }  $\equiv M \gg= (\lambda x. \text{do}\{N\})$     do{ $M; N$ }  $\equiv M \gg= (\lambda \_. \text{do}\{N\})$     do{ $M$ }  $\equiv M$ 
```

possibly trading semicolons and brackets for the conventional Haskell *layout*. For instance, the above example is rendered as

```
do
  x  $\leftarrow$  newVar 0
  fork (writeVar x 42)
  readVar x
```

## 2.2 Operational Semantics

We present the operational semantics of *OCTM* in terms of an abstract machine whose states are triples  $\langle P; \Theta, \Delta \rangle$  formed by

- thread family (process)  $P$ ;
- heap memory  $\Theta : \text{Loc} \rightarrow \text{Term}$ ;
- distributed working memory  $\Delta : \text{Loc} \rightarrow \text{Term} \times \text{TrName}$

where  $\text{Term}$  denotes the set of *OCTM* terms (cf. Figure 1) and  $\text{TrName}$  denotes the set of names used by the machine to identify active transactions.

We shall denote the set of all possible states as  $\text{State}$ .

---

<sup>2</sup>We treat the application of monadic combinators (e.g. `return`) as values in the line of similar works [7].

**Threads** Threads are the smaller unit of execution the machine scheduler operates on; they execute OCTM terms and do not have any private transactional memory.

Threads are given unique identifiers (ranged over by  $t$  or variations thereof) and, whenever they take part to some transaction, the transaction identifier (ranged over  $k, j$  or variations thereof). Threads of the former case are represented by  $(M)_t$  where  $M$  is the term being evaluated and the subscript  $t$  is the thread identifier. Threads of the latter case have two forms:  $(M \triangleright M'; N)_{t,k}$ , called and  $(M \triangleright M')_{t,k}$  where:

- $M$  is the term being evaluated inside the transaction  $k$ ;
- $M'$  is the term being evaluated as *compensation* in case  $k$  is aborted;
- $N$  is the term being evaluated as *continuation* after  $k$  commits or aborts.

Threads with a continuation are called *primary participants* (*to transaction  $k$* ), while threads without continuation are the *secondary participants*. The former group includes all and only the threads that started a transaction (i.e. those evaluated in an `atomic`), while the latter group encompasses threads forked inside a transaction and threads forced to join a transaction (from outside a transactional context) because of memory interactions. While threads of both groups can force a transaction to abort or restart, only primary participants can vote for its commit and hence pass the transaction result to the continuation.

We shall present thread families using the evocative CCS-like parallel operator  $\parallel$  (cf. Figure 2) which is commutative and associative. Notice that this operator is well-defined only on operands whose thread identifiers are distinct. The notation is extended to thread families with  $\mathbf{0}$  denoting the empty family.

**Memory** The memory is divided in the heap  $\Theta$  and in a distributed working memory  $\Delta$ . As for traditional closed (acid) transactions (e.g. [7]), operations inside a transaction are evaluated against  $\Delta$  and effects are propagated to  $\Theta$  only on commits. When a thread inside a transaction  $k$  accesses a location outside  $\Delta$  the location is *claimed for  $k$*  and remains claimed for the rest of  $k$  execution. Threads inside a transaction can interact only with locations claimed by their transaction. To this end, threads outside any transaction can join an existing one and different active transactions can be merged to share their claimed locations.

We shall denote the pair  $\langle \Theta, \Delta \rangle$  by  $\Sigma$  and reference to each projected component by a subscript e.g.  $\Sigma_\Theta$  for the heap. When describing updates to the state  $\Sigma$ , we adopt the convention that  $\Sigma'$  has to be intended as equal to  $\Sigma$  except if stated otherwise, i.e. by statements like  $\Sigma'_\Theta = \Sigma_\Theta[r \mapsto M]$ .

Formally, updates to location content are defined on  $\Theta$  and  $\Delta$  as follows:

$$\Theta[r \mapsto M](s) \triangleq \begin{cases} M & \text{if } r = s \\ \Theta(s) & \text{otherwise} \end{cases} \quad \Delta[r \mapsto (M, k)](s) \triangleq \begin{cases} (M, k) & \text{if } r = s \\ \Delta(s) & \text{otherwise} \end{cases}$$

for any  $r, s \in \text{Loc}$ ,  $M \in \text{Term}$  and  $k \in \text{TrName}$ . Likewise, updates on transaction names are defined on  $\Sigma$  and  $\Delta$  as follows:

$$\Sigma[k \mapsto j] \triangleq (\Theta, \Delta[k \mapsto j]) \quad (\Delta[k \mapsto j])(r) \triangleq \begin{cases} \Delta(r) & \text{if } \Delta(r) = (M, l), l \neq k \\ (M, j) & \text{if } \Delta(r) = (M, k) \end{cases}$$

for any  $r \in \text{Loc}$ ,  $M \in \text{Term}$  and  $k, j \in \text{TrName}$ . Note that  $j$  may occur in  $\Delta$  resulting in the fusion of the transactions denoted by  $k$  and  $j$  respectively. Finally,  $\emptyset$  denotes the empty memory (i.e. the completely undefined partial function).

Thread	$T_t ::= (\langle M \rangle)_t \mid (\langle M \triangleright M'; N \rangle)_{t,k} \mid (\langle M \triangleright M' \rangle)_{t,k}$
Thread family	$P ::= T_{t_1} \parallel \dots \parallel T_{t_n} \quad \forall i, j \ t_i \neq t_j$
Expressions	$\mathbb{E} ::= [-] \mid \mathbb{E} \gg= M$
Processes	$\mathbb{P}_t ::= (\langle \mathbb{E} \rangle)_t$
Transactions	$\mathbb{T}_{t,k} ::= (\langle \mathbb{E} \triangleright M; N \rangle)_{t,k} \mid (\langle \mathbb{E} \triangleright M \rangle)_{t,k}$

Figure 2: Threads and evaluation contexts.

$$\begin{array}{c}
 \frac{M \not\equiv V \quad \mathcal{V}[M] = V}{M \rightarrow V} \text{ (EVAL)} \quad \frac{\text{return } M \gg= N \rightarrow NM}{\text{return } M \gg= N \rightarrow NM} \text{ (BINDRETURN)} \\
 \hline
 \frac{}{\text{retry } \gg= M \rightarrow \text{retry}} \text{ (BINDRETRY)} \quad \frac{}{\text{abort } N \gg= M \rightarrow \text{abort } N} \text{ (BINDABORT)}
 \end{array}$$

Figure 3: *OCTM* semantics: rules for term evaluation.

**Behaviour** Evaluation contexts are shown in Figure 2 and the transition relations are presented in Figures 3, 4, 5. The first (cf. Figures 3) is defined on terms only and models pure computations.

In particular, rule (EVAL) allows a term  $M$  that is not a value to be evaluated by an auxiliary (partial) function,  $\mathcal{V}[M]$  yielding the value  $V$  of  $M$  whereas the other three rules define the semantic of the monadic bind. The transition relation modelling pure computations can be thought as accessory to the remaining two for these model transitions between the states of the machine under definition.

Derivation rules in Figure 4 characterize the execution of pure (effect-free) terms, forks and memory operations both inside, and outside of some transaction; Derivation rules in Figure 5 characterize auxiliary operations for transaction management (e.g. creation) and their coordination (e.g. distributed commits). Note that there are no derivation rules for `retry`. In fact, the meaning of `retry` is to inform the machine that choices made by the scheduler led to a state from which the program cannot proceed. From an implementation perspective this translates in the transaction being re-executed from the beginning (or a suitable check-point) following a different scheduling of its operations.

We shall describe now a representative subset of the derivation rules from Figures 4 and 5.

Reading a location falls into four cases depending on the location being claimed (i.e. occurring in  $\Delta$ ) and the reader being part of a transaction. The rule (READP) characterize the reading of an unclaimed location from outside any transaction; the read is performed as expected leaving it unclaimed. Rule (READT) describes the reading of an unclaimed location  $r$  by a thread belonging to some transaction  $k$ ; the side effect of the reading is  $r$  being claimed for  $k$ . Rules (READMERGE) and (READJOIN) cover the cases of readings against claimed locations. In the first scenario, the reading thread belongs to a transaction resulting in the two being merged, which is expressed by renaming its transaction via a substitution. In the remaining scenario, the reading thread does not belong to any transaction and hence joins the transaction  $k$  which claimed the location. The newly created participant does not have any continuation since the whole term is set to be executed inside  $k$ ; any other choice for splitting the term singling out a compensation would impose an artificial synchronization with the transaction commit. For a counter example, consider executing only the read operation inside the transaction and delaying everything after the commit; then concurrency will be clearly reduced. Because of the same reasoning, the whole term  $M$  is taken as the compensation of the participant.

Transactions are created by rule (ATOMIC); threads participating in a transaction are non-deterministically interleaved with other threads. The stronger requirement of isolation is offered

$$\begin{array}{c}
\frac{M \rightarrow N}{\langle \mathbb{P}_t[M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{P}_t[N] \parallel P; \Sigma \rangle} \text{ (TERMP)} \quad \frac{M \rightarrow N}{\langle \mathbb{T}_{t,k}[M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,k}[N] \parallel P; \Sigma \rangle} \text{ (TERM T)} \\
\frac{t' \notin \text{threads}(P) \quad t \neq t'}{\langle \mathbb{P}_t[\text{fork } M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{P}_t[\text{return } t'] \parallel \langle M \rangle_{t'} \parallel P; \Sigma \rangle} \text{ (FORKP)} \\
\frac{t' \notin \text{threads}(P) \quad t \neq t'}{\langle \mathbb{T}_{t,k}[\text{fork } M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,k}[\text{return } t'] \parallel \langle M \triangleright \text{return} \rangle_{t',k} \parallel P; \Sigma \rangle} \text{ (FORKT)} \\
\text{threads}(T_{t_1} \parallel \dots \parallel T_{t_n}) \triangleq \{t_1, \dots, t_n\} \\
\frac{r \notin \text{dom}(\Sigma_\Theta) \cup \text{dom}(\Sigma_\Delta) \quad \Sigma'_\Theta = \Sigma_\Theta[r \mapsto M]}{\langle \mathbb{P}_t[\text{newVar } M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{P}_t[\text{return } r] \parallel P; \Sigma' \rangle} \text{ (NEWP)} \\
\frac{r \notin \text{dom}(\Sigma_\Theta) \cup \text{dom}(\Sigma_\Delta) \quad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M, k)]}{\langle \mathbb{T}_{t,k}[\text{newVar } M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,k}[\text{return } r] \parallel P; \Sigma' \rangle} \text{ (NEWT)} \\
\frac{r \notin \text{dom}(\Sigma_\Delta) \quad \Sigma_\Theta(r) = M}{\langle \mathbb{P}_t[\text{readVar } r] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{P}_t[\text{return } M] \parallel P; \Sigma \rangle} \text{ (READP)} \\
\frac{r \notin \text{dom}(\Sigma_\Delta) \quad \Sigma_\Theta(r) = M \quad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M, k)]}{\langle \mathbb{T}_{t,k}[\text{readVar } r] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,k}[\text{return } M] \parallel P; \Sigma' \rangle} \text{ (READT)} \\
\frac{M = \mathbb{E}[\text{readVar } r] \quad \Sigma_\Delta(r) = (M', k)}{\langle \langle M \rangle_t \parallel P; \Sigma \rangle \rightarrow \langle \langle \mathbb{E}[\text{return } M'] \triangleright \lambda.M \rangle_{t,k} \parallel P; \Sigma \rangle} \text{ (READJOIN)} \\
\frac{\Sigma_\Delta(r) = (M, j) \quad \Sigma' = \Sigma[k \mapsto j]}{\langle \mathbb{T}_{t,k}[\text{readVar } r] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,j}[\text{return } M] \parallel P[k \mapsto j]; \Sigma' \rangle} \text{ (READMERGE)} \\
\frac{r \notin \text{dom}(\Sigma_\Delta) \quad \Sigma_\Theta(r) = N \quad \Sigma'_\Theta = \Sigma_\Theta[r \mapsto M]}{\langle \mathbb{P}_t[\text{writeVar } r M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{P}_t[\text{return } ()] \parallel P; \Sigma' \rangle} \text{ (WRITEP)} \\
\frac{r \notin \text{dom}(\Sigma_\Delta) \quad \Sigma_\Theta(r) = N \quad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M, k)]}{\langle \mathbb{T}_{t,k}[\text{writeVar } r M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,k}[\text{return } ()] \parallel P; \Sigma' \rangle} \text{ (WRITET)} \\
\frac{M = \mathbb{E}[\text{writeVar } r M'] \quad \Sigma_\Delta(r) = (M'', k) \quad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M', k)]}{\langle \langle M \rangle_t \parallel P; \Sigma \rangle \rightarrow \langle \langle \mathbb{E}[\text{return } ()] \triangleright \lambda.M \rangle_{t,k} \parallel P; \Sigma' \rangle} \text{ (WRITEJOIN)} \\
\frac{\Sigma_\Delta(r) = (N, j) \quad \Sigma' = \Sigma[k \mapsto j] \quad \Sigma'_\Delta = \Sigma_\Delta[k \mapsto (M, j)]}{\langle \mathbb{T}_{t,k}[\text{writeVar } r M] \parallel P; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,j}[\text{return } ()] \parallel P[k \mapsto j]; \Sigma' \rangle} \text{ (WRITEMERGE)}
\end{array}$$

Figure 4: *OCTM* semantics: rules for  $\rightarrow$ .

by rules (ISOLATEDP) and (ISOLATEDT), whose premises forbid thread or transaction creation.

Committing or aborting a transaction require a synchronization of its participants. In particular, an abort can be read as a participant vetoing the outcome of the transaction; this corresponds to (RAISEABORT1) and (RAISEABORT2). The information is then propagated by (ABBROADCAST) and (TRIGNORE) to any other participant to the transaction being aborted; these participants abort performing a transition described by either (SIGABORT1) or (SIGABORT2).

### 3 Examples

In this section we provide some short examples to illustrate the use of *OCTM* and how standard STM behaviour can be recovered in *OCTM* thanks to the *isolated* construct. In Section 4.2 we will give an extended example by providing a translation of *TCCS*<sup>m</sup> into *OCTM*.

$$\begin{array}{c}
\frac{k \notin \text{transactions}(P)}{\langle (\text{atomic } M \text{ } N \gg= N')_t \parallel P; \Sigma \rangle \xrightarrow{\text{new}_k} \langle (\langle M \triangleright N; N' \rangle_{t,k} \parallel P; \Sigma \rangle} \text{(ATOMIC)} \\
\frac{\langle (\langle M \rangle_t; \Sigma \rangle \xrightarrow{*} \langle (\text{return } N)_t; \Sigma' \rangle}{\langle \mathbb{P}_t[\text{isolated } M]; \Sigma \rangle \rightarrow \langle \mathbb{P}_t[\text{return } N]; \Sigma' \rangle} \text{(ISOLATEDP)} \\
\frac{\text{op} \in \{\text{abort, return}\} \quad \langle (\langle M \triangleright \text{return} \rangle_{t,k}; \Sigma \rangle \xrightarrow{*} \langle (\text{op } N \triangleright \text{return})_{t,k}; \Sigma' \rangle}{\langle \mathbb{T}_{t,k}[\text{isolated } M]; \Sigma \rangle \rightarrow \langle \mathbb{T}_{t,k}[\text{op } N]; \Sigma' \rangle} \text{(ISOLATEDT)} \\
\frac{\Sigma'_\Delta = \text{clean}(k, \Sigma_\Delta)}{\langle (\text{abort } M \triangleright N; N')_{t,k}; \Sigma \rangle \xrightarrow{\text{ab}_k M} \langle (\langle N(M) \gg= N' \rangle_t; \Sigma' \rangle} \text{(RAISEABORT1)} \\
\frac{\Sigma'_\Delta = \text{clean}(k, \Sigma_\Delta)}{\langle (\text{abort } M \triangleright N)_{t,k}; \Sigma \rangle \xrightarrow{\text{ab}_k M} \langle (\langle N(M) \rangle_t; \Sigma' \rangle} \text{(RAISEABORT2)} \\
\frac{\Sigma'_\Delta = \text{clean}(k, \Sigma_\Delta)}{\langle (\langle M \triangleright N; N' \rangle_{t,k}; \Sigma \rangle \xrightarrow{\widehat{\text{ab}}_k M} \langle (\langle N(M) \gg= N' \rangle_t; \Sigma' \rangle} \text{(SIGABORT1)} \\
\frac{\Sigma'_\Delta = \text{clean}(k, \Sigma_\Delta)}{\langle (\langle M \triangleright N)_{t,k}; \Sigma \rangle \xrightarrow{\widehat{\text{ab}}_k M} \langle (\langle N(M) \rangle_t; \Sigma' \rangle} \text{(SIGABORT2)} \\
\frac{\langle P; \Sigma \rangle \xrightarrow{\text{ab}_k M} \langle P'; \Sigma' \rangle \quad \langle Q; \Sigma \rangle \xrightarrow{\widehat{\text{ab}}_k M} \langle Q'; \Sigma' \rangle}{\langle P \parallel Q; \Sigma \rangle \xrightarrow{\text{ab}_k M} \langle P' \parallel Q'; \Sigma' \rangle} \text{(ABBROADCAST)} \\
\frac{\Sigma'_\Theta = \text{commit}(k, \Sigma_\Theta, \Sigma_\Delta) \quad \Sigma'_\Delta = \text{clean}(k, \Sigma_\Delta)}{\langle (\text{return } M \triangleright N; N')_{t,k}; \Sigma \rangle \xrightarrow{\text{co}_k} \langle (\text{return } M \gg= N')_t; \Sigma' \rangle} \text{(COMMIT1)} \\
\frac{\Sigma'_\Theta = \text{commit}(k, \Sigma_\Theta, \Sigma_\Delta) \quad \Sigma'_\Delta = \text{clean}(k, \Sigma_\Delta)}{\langle (\langle M \triangleright N \rangle_{t,k}; \Sigma \rangle \xrightarrow{\text{co}_k} \langle (\langle M \rangle_t; \Sigma' \rangle} \text{(COMMIT2)} \\
\frac{\langle P; \Sigma \rangle \xrightarrow{\text{co}_k} \langle P'; \Sigma' \rangle \quad \langle Q; \Sigma \rangle \xrightarrow{\text{co}_k} \langle Q'; \Sigma' \rangle}{\langle P \parallel Q; \Sigma \rangle \xrightarrow{\text{co}_k} \langle P' \parallel Q'; \Sigma' \rangle} \text{(CoBROADCAST)} \\
\frac{\langle P; \Sigma \rangle \xrightarrow{\beta} \langle P'; \Sigma' \rangle \quad \text{transactions}(\beta) \notin \text{transactions}(Q)}{\langle P \parallel Q; \Sigma \rangle \xrightarrow{\beta} \langle P' \parallel Q; \Sigma \rangle} \text{(TRIGNORE)} \\
\text{clean}(k, \Delta)(r) \triangleq \begin{cases} \perp & \text{if } \Delta(r) = (M, k) \\ \Delta(r) & \text{otherwise} \end{cases} \quad \text{commit}(k, \Theta, \Delta)(r) \triangleq \begin{cases} M & \text{if } \Delta(r) = (M, k) \\ \Theta(r) & \text{otherwise} \end{cases} \\
\text{transactions}(\langle M \rangle_t) \triangleq \emptyset \quad \text{transactions}(\langle M \triangleright M'; N \rangle_{t,k}) \triangleq \{k\} \quad \text{transactions}(\langle M \triangleright N \rangle_{t,k}) \triangleq \{k\} \\
\text{transactions}(T_{t_1} \parallel \dots \parallel T_{t_n}) \triangleq \bigcup_{1 \leq i \leq n} \text{transactions}(T_{t_i})
\end{array}$$

Figure 5: *OCTM* semantics: rules for  $\xrightarrow{\beta}$ .

### 3.1 MVars

One of the basic constructs offered by Concurrent Haskell are *MVars* [10] i.e. mutable locations that are either empty or contain a value of the given type parameter. Interaction with these structures is based on two fundamental operations: `putMVar` which fills an *MVar* if it is empty and blocks otherwise, and `takeMVar` which empties an *MVar* if it is full and blocks otherwise. In [7] *MVars* are implemented on top of *TVars* (i.e. STM Haskell transactional locations).

Following [7] an `MVar` of type `a` is implemented on top of a `OTVar` (our transactional locations i.e. any  $r \in \text{Loc}$ ) holding a value of type `Maybe a`; this is a type that is either an empty value (`Nothing`) or actually holds a value of type `a` (e.g. `Just 42`). Thus, the definition of the type `MVar a` is the following:

```
type MVar a = OTVar (Maybe a)
```

and its two constructors for creating an empty and a full location are:

```
newEmptyMVar = newVar Nothing
newMVar x = newVar (Just x)
```

The definition of the two basic operations is precisely the same appearing in [7] except for the added `isolated` construct for enforcing isolation.

```
takeMVar v = isolated do
  v ← readVar v
  case v of
    Nothing → retry
    Just x → do
      writeVar x Nothing
      return x

putMVar v y = isolated do
  v ← readVar v
  case v of
    Nothing → writeVar y Nothing
    Just _ → retry
```

### 3.2 Transactional RPC

`MVars` can be used as simple directional channels with `takeMVar` and `putMVar` as receive and send. Then a bidirectional channel for a *remote procedure call* is easily implemented using a pair of `MVars`

```
type RPC a b = (MVar (CorId, a), MVar (CorId, b))
```

where `a` and `b` are the types of the request and response exchanged and `CorId` is a suitable type providing a correlation identifier for relating a request to its response.

Before we introduce the skeleton and stub let us define a conditional variation of the `takeMVar` accepting a boolean predicate `p` and such that it empties the given `MVar v` only if the contained value satisfies `p` and blocks (issue a `retry`) otherwise.

```
takeMVarIf p v = isolated do
  v ← readVar v
  case v of
    Nothing → retry
    Just x → do
      if p x then
        writeVar x Nothing >> return x
      else
        retry
```

The conditional version of `takeMVar` allows us to take a response only if we know its correlation identifier and hence the call is simply:

```
rpcCall (req, res) data = do
  c ← newCorrelationId
  putMVar req (c, data)
  r ← takeMVarIf (c == fst) res
  return (snd r)
```

where `fst` and `snd` are the first and second projections respectively. Symmetrically, to provide the `rpc` we just need to take a request from the `MVar` `req` and put its response in `res` using the same correlation identifier:

```
rpcServe (req, res) data = do
  q ← takeMVar req
  a ← doSomething (snd q)
  putMVar res (fst q, a)
```

If any of the two parties happens to be partaking a transaction the `rpc` results in the other joining the transaction effectively rendering the `rpc` transactional.

The above example is quite simplified (e.g. requests could have been handled by a buffer, and the structure of `(req, req)` should be hidden to the user) but serves the purpose of illustrating the difference between *OCTM* and *STM*. handled by a buffer, and the structure of `(req, req)` should be hidden to the user) but serves the purpose of illustrating the difference between *OCTM* and *STM*. In fact, the above implementation allows the call to happen inside a transaction without resulting into a lock as in the case of *STM* since isolation will prevent the serving thread to join and provide a response.

## 4 Expressiveness of *OCTM*

In order to assess the expressive power of *OCTM*, in this Section we prove that it can be used to implement *TCCS<sup>m</sup>*, a formal model for open transactions [11]. We proceed as follow: first, in Subsection 4.1 we recall *TCCS<sup>m</sup>*; then, the translation of *TCCS<sup>m</sup>* processes into *OCTM* states is defined in Subsection 4.2; this translation is proved to be correct in Subsection 4.3.

### 4.1 *TCCS<sup>m</sup>*: CCS with open transactions

*TCCS<sup>m</sup>* [11] is a CCS-like calculus with open flat can synchronize even when belonging to different transactions, which in turn are joined into a distributed one. We refer to [11] for a detailed description of *TCCS<sup>m</sup>*. transactions: processes can synchronize even when belonging to different transactions, which in turn are joined into a distributed one. We refer to [11] for a detailed description of *TCCS<sup>m</sup>*.

The syntax of *TCCS<sup>m</sup>* is defined by the following grammar

$$P ::= \sum_{i=1}^n \alpha_i.P_i \mid \prod_{i=0}^m P_i \mid P \setminus L \mid X \mid \mu X.P \mid \langle P_1 \blacktriangleright P_2 \rangle \mid \langle\langle P_1 \triangleright_k P_2 \rangle\rangle \mid \text{co}.P \quad (1)$$

where  $\alpha_i ::= a \mid \bar{a} \mid \tau$ ,  $a$  ranges over a given set of visible actions  $A$ ,  $L$  over subsets of  $A$  and the bijection  $(\bar{\cdot}) : A \rightarrow A$  maps every action to its *coaction* as usual. The calculus extends CCS with three constructs which represent *inactive* transactions, *active* transactions and *commit* actions respectively. Transactions such as  $\langle\langle P_1 \triangleright_k P_2 \rangle\rangle$  are formed by two processes with the former being executed atomically and the latter being executed whenever the transaction is aborted, i.e. as a *compensation*. Terms denoting active transactions expose also a name ( $k$  in the previous example) which is used to track transaction fusions. For instance, consider the process denoted by  $\langle\langle P_1 \triangleright_j P_2 \rangle\rangle \mid \langle\langle Q_1 \triangleright_k Q_2 \rangle\rangle$  where  $P_1$  and  $Q_1$  synchronize on some  $a \in A$ ; the result of this synchronization is the fusion of the transactions  $j$  and  $k$  i.e.  $\langle\langle P'_1 \triangleright_l P_2 \rangle\rangle \mid \langle\langle Q'_1 \triangleright_l Q_2 \rangle\rangle$ . The fusion makes explicit the dependency between  $j$  and  $k$  introduced by the synchronization and ties them to agree on commits. In this sense,  $P'_1$  and  $Q'_1$  are participants of a *distributed transaction* [6].

As in [11] we restrict ourselves to well-formed terms. Intuitively, a term is well-formed if active transactions occur only at the top-level and commit actions occur only in a transaction (active

$$\begin{array}{c}
\frac{\varsigma \vdash P : \mathbf{p}}{\varsigma \vdash P : \tau} \quad \frac{\varsigma \vdash P : \mathbf{p}}{\varsigma \vdash \text{co}.P : \mathbf{c}} \quad \frac{\varsigma \vdash P : \tau}{\varsigma \vdash P \setminus L : \tau} \\
\frac{}{\varsigma \vdash X : \varsigma(X)} \quad \frac{\varsigma[X : \mathbf{p}] \vdash P : \mathbf{p}}{\varsigma \vdash \mu X.P : \mathbf{p}} \quad \frac{\varsigma[X : \mathbf{c}] \vdash P : \mathbf{c}}{\varsigma \vdash \mu X.P : \mathbf{c}} \quad \frac{\forall i \varsigma \vdash P_i : \tau}{\varsigma \vdash \prod P_i : \tau} \\
\frac{\forall i \varsigma \vdash P_i : \mathbf{p}}{\varsigma \vdash \sum \alpha_i.P_i : \mathbf{p}} \quad \frac{\forall i \varsigma \vdash \alpha_i.P_i : \mathbf{c}}{\varsigma \vdash \sum \alpha_i.P_i : \mathbf{c}} \quad \frac{\varsigma \vdash P : \mathbf{c}}{\varsigma \vdash \langle\langle P \triangleright_k Q \rangle\rangle : \mathbf{t}} \quad \frac{\varsigma \vdash P : \mathbf{c} \quad \varsigma \vdash Q : \mathbf{p}}{\varsigma \vdash \langle P \blacktriangleright Q \rangle : \mathbf{p}}
\end{array}$$

Figure 6: Simple types for  $TCCS^m$ .

or inactive). To this end we introduce a *type system* for  $TCCS^m$ , whose rules are in Figure 6. Terms that cannot occur inside a transaction have type  $\mathbf{t}$ , terms that cannot occur outside a transaction have type  $\mathbf{c}$ , and terms without such restrictions have type  $\mathbf{p}$ ;  $\tau$  ranges over types.

**Definition 1** (Well-formed  $TCCS^m$  terms). *A  $TCCS^m$  term  $P$ , described by the grammar in (1), is said to be well-formed if, and only if,  $\emptyset \vdash P : \mathbf{t}$ . Well-formed terms form the set  $\text{Proc}$ .*

The operational semantics of well-formed  $TCCS^m$  terms is given by the SOS in Figure 7 (see [11] for further details). The reduction semantics is given as a binary relation  $\rightarrow$  defined by

$$P \rightarrow Q \stackrel{\Delta}{\iff} P \xrightarrow{\tau} Q \vee P \xrightarrow{\beta} Q \vee P \xrightarrow{k(\tau)} Q.$$

The first case is a synchronization between pure CCS processes. The second case corresponds to creation of new transactions and distributed commit or abort ( $\beta \in \{\text{newk}, \text{cok}, \text{abk}\}$ ). The third case corresponds to synchronizations of processes inside a named (and possibly distributed) transaction. Notice that by (TSYNC) transaction fusion is *driven by communication* and that by (TSUM) any pure CCS process can join and interact with a transaction.

## 4.2 Encoding $TCCS^m$ in $OCTM$

In this section we define the translation from  $TCCS^m$  processes to  $OCTM$  states. To this end, we have to implement transactions and CCS-like synchronizations using shared transactional variables and the `atomic` and `isolated` operators.

Synchronization is implemented by means of shared transactional variables, one for each channel, that take values of type `ChState` (cf. Figure 9); this type has four constructors: one for each of the three messages of the communication protocol below plus a “nothing” one providing the default value. Let  $t_1$  and  $t_2$  be the identifiers of two threads simulating  $a.P$  and  $\bar{a}.Q$  respectively. The protocol is composed by the following four steps:

1.  $t_1$  checks whether the channel is free and writes on the transactional variable modelling the channel  $a$  a nonce tagged with the constructor `M1`;
2.  $t_2$  reads the variable for  $a$  and accepts the synchronization offered by the challenge (`M1 np`) adding a fresh nonce to it and writing back (`M2 np nq`);
3.  $t_1$  reads the answer to its challenge and acknowledges the synchronization writing back the nonce it read tagged with the constructor `M3`;
4.  $t_2$  reads the acknowledgement and frees the channel.

$$\begin{array}{c}
\frac{\sum \alpha_i.P_i \xrightarrow{\alpha_i} \varepsilon P_i}{\sum \alpha_i.P_i} \text{ (SUM)} \quad \frac{P \xrightarrow{a} \varepsilon P' \quad Q \xrightarrow{\bar{a}} \varepsilon Q'}{P|Q \xrightarrow{\tau} \varepsilon P'|Q'} \text{ (SYNC)} \quad \frac{}{\mu X.P \xrightarrow{\tau} \varepsilon P[\mu X.P/X]} \text{ (REC)} \\
\frac{P \xrightarrow{\alpha} \sigma P' \quad \text{img}(\sigma) \cap \text{tn}(Q) = \emptyset}{P|Q \xrightarrow{\alpha} \sigma P'|Q[\sigma]} \text{ (PARL)} \quad \frac{\tau \neq \alpha_j}{\sum \alpha_i.P_i \xrightarrow{k(\alpha_j)} \varepsilon \mapsto k \langle\langle P_j | \text{co} \triangleright_k \sum \alpha_i.P_i \rangle\rangle} \text{ (T SUM)} \\
\frac{P \xrightarrow{\alpha} \varepsilon P' \quad \tau \neq \alpha \quad l \neq k}{\langle\langle P \triangleright_l Q \rangle\rangle \xrightarrow{k(\alpha)} \varepsilon \mapsto k \langle\langle P' \triangleright_k Q \rangle\rangle} \text{ (TA CT)} \quad \frac{P \xrightarrow{k(a)} \varepsilon \mapsto k P' \quad Q \xrightarrow{k(\bar{a})} \varepsilon \mapsto k Q'}{P|Q \xrightarrow{k(\tau)} \varepsilon \mapsto k P'[j \mapsto k]|Q'[i \mapsto k]} \text{ (T SYNC)} \\
\frac{P \xrightarrow{\alpha} \sigma P' \quad \alpha \notin L}{P \setminus L \xrightarrow{\alpha} \sigma P' \setminus L} \text{ (RES)} \quad \frac{P \xrightarrow{\tau} \varepsilon P'}{\langle\langle P \triangleright_k Q \rangle\rangle \xrightarrow{\tau} \varepsilon \langle\langle P' \triangleright_k Q \rangle\rangle} \text{ (TT AU)} \quad \frac{}{\langle\langle P \triangleright_k Q \rangle\rangle \xrightarrow{\text{abk}} Q} \text{ (TAB)} \\
\frac{P \xrightarrow{\beta} P'}{P \setminus L \xrightarrow{\beta} P' \setminus L} \text{ (TRES)} \quad \frac{k \text{ fresh}}{\langle P \blacktriangleright Q \rangle \xrightarrow{\text{newk}} \langle\langle P \triangleright_k Q \rangle\rangle} \text{ (TNEW)} \quad \frac{\exists i P_i = \text{co}.P'_i}{\langle\langle \prod P_i \triangleright_k Q \rangle\rangle \xrightarrow{\text{cok}} \Psi_{id}(P)} \text{ (TCO)} \\
\frac{P \xrightarrow{\beta} P' \quad Q \xrightarrow{\beta} Q' \quad \beta \neq \text{newk}}{P|Q \xrightarrow{\beta} P'|Q'} \text{ (TB1)} \quad \frac{P \xrightarrow{\beta} P' \quad \text{tn}(\beta) \notin \text{tn}(Q)}{P|Q \xrightarrow{\beta} P'|Q} \text{ (TB2)} \\
\Psi_\sigma(P) \triangleq \begin{cases} Q & \text{if } P = \text{co}.Q \\ \Psi_\sigma(Q) \setminus L & \text{if } P = Q \setminus L \\ \sum \alpha_i.\Psi_\sigma(P_i) & \text{if } P = \sum \alpha_i.P_i \\ \prod \Psi_\sigma(P_i) & \text{if } P = \prod P_i \\ \mu X.\Psi_{\sigma[P/X]}(Q) & \text{if } P = \mu X.Q \\ P[\sigma] & \text{otherwise} \end{cases} \\
\text{tn}(P) \triangleq \begin{cases} \{k\} & \text{if } P = \langle\langle P \triangleright_k Q \rangle\rangle \\ \bigcup \text{tn}(P_i) & \text{if } P = \prod P_i \\ \emptyset & \text{otherwise} \end{cases} \\
\text{tn}(\beta) \triangleq \begin{cases} k & \text{if } \beta = \text{newk} \\ k & \text{if } \beta = \text{abk} \\ k & \text{if } \beta = \text{cok} \end{cases}
\end{array}$$

Figure 7:  $TCCS^m$  operational semantics.

Each step has to be executed in isolation with respect to the interactions with the shared transactional variable  $a$ .

Nonces are meant to correlate the steps only and hence can be easily implemented in  $OCTM$  by pairing thread identifiers with counter *a la* logical clock. If at any step a thread finds the channel in an unexpected state it means that the chosen scheduling has led to a state incoherent with respect to the above protocol; hence the thread executes a *retry*. This tells the scheduler to try another execution order; by fairness, we eventually find a scheduling such that the two processes do synchronize on  $a$  and these are the only executions leading to  $P | Q$ . The protocol is illustrated in Figure 8. If the synchronizing parties are involved in distinct transactions these are fused as a side effect of the interaction via the shared variable.

A choice like  $\sum_{i=1}^m \alpha_i.P_i$  can be seen as a race of threads  $t_1, \dots, t_m$ , each simulating a branch, to acquire a boolean transactional variable  $l$  (private to the group). Each  $t_i$  proceeds as follows. First, it checks  $l$  and if it is set, it returns void and terminates (another thread has already acquired it); otherwise it tries to set it while carrying out  $\alpha_i$ , i.e. right before executing its last step of the communication protocol. If the variable is acquired by another thread while  $t_i$  is finalizing  $\alpha_i$  then  $t_i$  issues a *retry* to retract any effect of  $\alpha_i$ . The  $OCTM$  code implementing this protocol is shown in Figure 9.

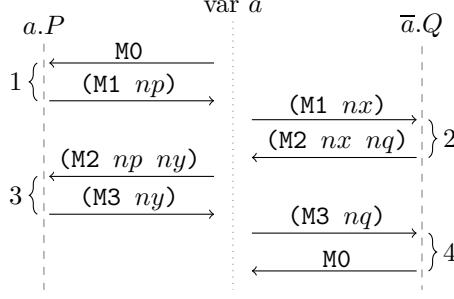


Figure 8: Implementing  $TCCS^m$  synchronization.

**Encoding of  $TCCS^m$**  We can now define the encoding  $\eta : \text{Proc} \rightarrow \text{State}$ , mapping well-formed  $TCCS^m$  terms to states of the  $OCTM$  abstract machine. Intuitively, a process  $P \equiv \prod_{i=1}^m P_i$  is mapped into a state with a thread for each  $P_i$  and a variable for each channel in  $P$ . Clearly a state of this form can be generated by a single  $OCTM$  term which allocates all variables and forks the  $m$  threads; we have preferred to map  $TCCS^m$  terms to  $OCTM$  states instead of  $OCTM$  term for sake of simplicity.

The map  $\eta$  is defined by recursion along the derivation of  $\emptyset \vdash P : t$  and the number  $m$  of parallel components in  $P \equiv \prod_{i=1}^m P_i$ . This is handled by the auxiliary encoding  $\varsigma : \text{Proc} \times \text{Heap} \rightarrow \text{State}$  (up to choice of fresh names) whose second argument is used to track memory allocations. The base case is given by  $m = 0$  and yields a state with no threads i.e.  $\langle \mathbf{0}, \Theta, \emptyset \rangle$ . The recursive step is divided in three subcases depending on the structure and type of  $P_1$  ( $m > 0$ ).

1. If  $\emptyset \vdash P_1 : c$  without top-level restrictions (i.e. for no  $Q$  and no  $L = \{a_1, \dots, a_{n+1}\}$  such that each  $a_i$  occurs in  $Q$  the process  $P_1$  is structurally equivalent to  $Q \setminus L$ ) then

$$\varsigma(\prod_{i=1}^{m+1} P_i, \Theta) \triangleq \langle [\varrho(P_1)]_{t_1} \parallel S; \Sigma \rangle$$

where  $\langle S; \Sigma \rangle = \varsigma(\prod_{j=1}^{m-1} P_{j+1}, \Theta)$  is the translation of the rest of  $P$  and  $t_1$  is unique w.r.t.  $S$  (i.e.  $t_1 \notin \text{threads}(S)$ ). By hypothesis  $P_1$  does not contain any top-level active transaction or parallel composition and hence can be translated directly into a  $OCTM$ -term

by means of the encoding  $\varrho$  (cf. Figure 10) –  $\varrho(P)$  contain a free variable for each unrestricted channel occurring in  $P$ .

2. If  $P_1$  has a top-level restriction (i.e.  $P_1 \equiv Q \setminus \{a_1, \dots, a_{n+1}\}$ ) then

$$\varsigma(\prod_{i=1}^{m+1} P_i, \Theta) \triangleq \langle S_1[r_1/a_1, \dots, r_{n+1}/a_{n+1}] \parallel S_2; \Theta_2[r_1, \dots, r_{n+1} \mapsto \text{M0}], \emptyset \rangle$$

where  $\langle S_1; \Theta_1, \emptyset \rangle = \varsigma(Q, \Theta)$  and  $\langle S_2; \Theta_2, \emptyset \rangle = \varsigma(\prod_{j=1}^{m-1} P_{j+1}, \Theta_1)$  are the translation of the unrestricted process  $Q$  and the translation of the rest of  $P$  respectively, all threads have a unique identifier  $\text{threads}(S_1) \cap \text{threads}(S_2) = \emptyset$ , the heap is extended with  $n$  channel variables fresh ( $r_1, \dots, r_{n+1} \notin \text{dom}(\Theta_2)$ ) and known only to the translation of  $Q$ .

3. If  $P_1 \equiv \langle\langle Q_1 \triangleright_k Q_2 \rangle\rangle$  is an active transaction then

$$\begin{aligned} \varsigma(\prod_{i=1}^{m+1} P_i, \Theta) &\triangleq \langle S_{co} \parallel S_{ab} \parallel S_1[r_{co}/co] \parallel S_2; \Theta_2[r_l \mapsto \text{True}, r_{co} \mapsto \text{M0}], \emptyset \rangle \\ S_{co} &= \langle \text{recv } r_l \ r_{co} \triangleright \varrho(Q_1); \text{bang } (\text{recv } (\text{newVar } \text{True}) \ r_{co}) \rangle_{t_{co}, k} \\ S_{ab} &= \langle \text{abort } () \triangleright \text{return} \rangle_{t_{ab}, k} \end{aligned}$$

```

data Channel = OTVar ChState
data ChState = M1 Nonce | M2 Nonce Nonce | M3 Nonce | M0

tau l P = isolated do
  case (readVar l) of
    False → return ()
    True → chooseThis l >> P

chooseThis l = writeVar l False

eqOrRetry x y
| x == y = return ()
| otherwise = retry

bang x = fork x >> bang x

recv c l P = do
  nq ← newNonce
  isolated do
    case (readVar 1) of
      False → return ()
      True → do
        chooseThis l
        case (readVar c) of
          (M1 nx) → writeVar c (M2 nx nq)
          _ → retry
  isolated do
    case (readVar c) of
      (M3 ny) → eqOrRetry ny nq >> writeVar c M0 >> P
      _ → retry

send c l P = do
  np ← newNonce
  isolated do
    case (readVar l) of
      False → return ()
      True → do
        chooseThis l
        case (readVar c) of
          M0 → writeVar c (M1 np)
          _ → retry
  isolated do
    case (readVar c) of
      (M2 nx ny) → eqOrRetry nx np >> writeVar c (M3 ny) >> P
      _ → retry

```

Figure 9: Encoding channels and communication

$$\begin{aligned}
\varrho(\sum_{i=1}^m \alpha_i P_i) &\triangleq \text{do} & \varrho(\mu X.P) &\triangleq \text{let } X = \varrho(P) \text{ in} \\
l &\leftarrow \text{newVar True} & \varrho(P \blacktriangleright Q) &\triangleq \text{do} \\
\forall i \in \{1, \dots, m\} & & co &\leftarrow \text{newVar M0} \\
\text{fork } \xi(\alpha_i, l, P_i) & & \text{atomic p } \varrho(Q) \\
\varrho(\prod_{i=0}^m P_i) &\triangleq \text{do} & \text{bang psi} \\
\forall i \in \{0, \dots, m\} & & \text{where} \\
\text{fork } \varrho(P_i) & & p &= \text{do} \\
\varrho(P \setminus L) &\triangleq \text{do} & \varrho(P) \\
\forall c \in L & & \text{fork (abort ())} \\
c &\leftarrow \text{newVar M0} & \text{psi} \\
\varrho(P) & & \text{psi} &= \text{do} \\
\varrho(X) &\triangleq X & l &\leftarrow \text{newVar True} \\
\varrho(P) & & \text{recv } co \text{ } l \text{ return} \\
\varrho(co.P) &\triangleq \text{do} & \xi(\alpha_i, l, P_i) &\triangleq \begin{cases} \text{recv } \alpha_i \text{ } l \text{ } \varrho(P_i) & \text{if } \alpha_i = c \\ \text{send } \overline{\alpha_i} \text{ } l \text{ } \varrho(P_i) & \text{if } \alpha_i = \overline{c} \\ \text{tau } l \text{ } \varrho(P_i) & \text{if } \alpha_i = \tau \end{cases} \\
l &\leftarrow \text{newVar True} & & \\
\text{send } co \text{ } l \text{ } \varrho(P) & & &
\end{aligned}$$

Figure 10: Encoding  $TCCS^m$  terms of type  $\mathbf{c}$

where  $\langle S_1; \Theta_1, \emptyset \rangle = \varsigma(Q_1, \Theta)$ ,  $\langle S_2; \Theta_2, \emptyset \rangle = \varsigma(\prod_{j=1}^{m-1} P_{j+1}, \Theta_2)$  (like above), the thread  $S_{ab}$  is always ready to abort  $k$  as in (TAB) and  $S_{co}$  awaits on the private channel  $r_{co}$  a thread from  $S_1$  to reach a commit and, after its commit, collects all remaining synchronizations on  $r_{co}$  to emulate the effect of  $\Psi$  (cf. Figure 7). Finally, all threads have to be uniquely identified:  $\text{threads}(S_1) \cap \text{threads}(S_2) = \emptyset$  and  $t_{co}, t_{ab} \notin \text{threads}(S_1) \cup \text{threads}(S_2)$

**Remark 1.** *The third case of the definition above can be made more precise (at the cost of a longer definition) since the number of commits to be collected can be inferred from  $Q$  mimicking the definition of  $\Psi$ . This solution reduces the presence of dangling auxiliary processes and transaction fusions introduced by the cleaning process.*

Like  $\varrho$ ,  $\varsigma(P, \Theta)$  contains a free variable for each unrestricted channel in  $P$ . Finally, the encoding  $\eta$  is defined on each  $P \in \text{Proc}$  as:

$$\eta(P) \triangleq \langle S[r_1/a_1, \dots, r_n/a_n]; \Theta[r_1, \dots, r_n \mapsto \text{M0}], \emptyset \rangle$$

where  $\langle S; \Theta, \emptyset \rangle = \varsigma(P, \emptyset)$ ,  $\{r_1, \dots, r_n\} \subseteq \text{Loc}$ , and  $\{a_1, \dots, a_n\} \subseteq A$  is the set of channels occurring in  $P$ .

### 4.3 Adequacy of translation

In this section we prove that the translation  $\eta$  is adequate, in the sense that it preserves the observational behaviour of  $TCCS^m$  processes. More precisely, akin to [12], we define an appropriate notion of *star simulation*  $\mathcal{S}$  between well-formed  $TCCS^m$  processes and states of  $OCTM$ . The basic idea is that a single step of  $P$  is simulated by a sequence of reductions of  $\eta(P)$ , and  $\eta(P)$  does not exhibit behaviours which are not exhibited by  $P$ .

**Definition 2** (Star simulation). A relation  $\mathcal{S} \subseteq \text{Proc} \times \text{State}$  is a star simulation if for all  $(P, \langle S; \Sigma \rangle) \in \mathcal{S}$ :

1. for all  $Q$  such that  $P \xrightarrow{\tau} \sigma Q$  or  $P \xrightarrow{k(\tau)} \sigma Q$ , there exist  $S', \Sigma'$  such that  $\langle S; \Sigma \rangle \xrightarrow{*} \langle S'; \Sigma' \rangle$  and  $(Q, \langle S'; \Sigma' \rangle) \in \mathcal{S}$ ;
2. for all  $Q$  such that  $P \xrightarrow{\beta} Q$ , there exist  $S', \Sigma'$  s.t.  $\langle S; \Sigma \rangle \xrightarrow{\beta}^* \langle S'; \Sigma' \rangle$  and  $(Q, \langle S'; \Sigma' \rangle) \in \mathcal{S}$ .
3. for all  $S', \Sigma'$  such that  $\langle S; \Sigma \rangle \rightarrow \langle S'; \Sigma' \rangle$ , there exist  $Q, S'', \Sigma''$  such that  $(Q, \langle S''; \Sigma'' \rangle) \in \mathcal{S}$  and one of the following holds:
  - $P \xrightarrow{\tau} \sigma Q$  or  $P \xrightarrow{k(\tau)} \sigma Q$ , and  $\langle S'; \Sigma' \rangle \xrightarrow{*} \langle S''; \Sigma'' \rangle$
  - $P \xrightarrow{\beta} \epsilon Q$  and  $\langle S'; \Sigma' \rangle \xrightarrow{\beta}^* \langle S''; \Sigma'' \rangle$ .

where  $\beta$ -labels of the two transition relations are considered equivalent whenever are both commits or both aborts for the same transaction name. We say that  $P$  is star-simulated by  $\langle S; \Sigma \rangle$  if there exists a star-simulation  $\mathcal{S}$  such that  $(P, \langle S; \Sigma \rangle) \in \mathcal{S}$ . We denote by  $\approx$  the largest star simulation.

Another technical issue is that two equivalent  $TCCS^m$  processes can be translated to  $OCTM$  states which differ only on non-observable aspects, like name renamings, terminated threads, etc. To this end, we need to consider  $OCTM$  states up-to an equivalence relation  $\cong_t \subseteq \text{State} \times \text{State}$ , which we define next.

**Definition 3.** Two  $OCTM$  states are transaction-equivalent, written  $\langle S_1; \Sigma_1 \rangle \cong_t \langle S_2; \Sigma_2 \rangle$ , when they are equal up to:

- renaming of transaction and thread names;
- terminated threads, i.e. threads of one of the following forms:  $(\text{return } M)_t$ ,  $(\text{abort } M)_t$ ,  $(\text{return} \triangleright \text{return})_{t,k}$ ,  $(\text{abort} \triangleright \text{return})_{t,k}$ ,  $(\text{psi})_t$ ;
- threads blocked in synchronizations on co variables.

**Definition 4.** Let  $P \in \text{Proc}$  be a well-formed process and  $\langle S; \Sigma \rangle$  be a state.  $P$  is star simulated by  $\langle S; \Sigma \rangle$  up to  $\cong_t$  if  $(P, \langle S; \Sigma \rangle) \in \approx \circ \cong_t$ .

We are now ready to state our main adequacy result, which is a direct consequence of the two next technical lemmata.

**Lemma 1.** For all  $P, Q \in \text{Proc}$  the following hold true:

1. if  $P \xrightarrow{\tau} \sigma Q$  or  $P \xrightarrow{k(\tau)} \sigma Q$ , there exist  $S, \Sigma$  such that  $\eta(P) \xrightarrow{*} \langle S; \Sigma \rangle$  and  $\langle S; \Sigma \rangle \cong_t \eta(Q)$ ;
2. if  $P \xrightarrow{\beta} Q$ , there exist  $S, \Sigma$  such that  $\eta(P) \xrightarrow{\beta}^* \langle S; \Sigma \rangle$  and  $\langle S; \Sigma \rangle \cong_t \eta(Q)$ .

*Proof.* See Appendix A. □

**Lemma 2.** For  $P \in \text{Proc}$ , for all  $S, \Sigma$ , if  $\eta(P) \rightarrow \langle S; \Sigma \rangle$  then there exist  $Q, S', \Sigma'$  such that  $\langle S'; \Sigma' \rangle \cong_t \eta(Q)$  and one of the following holds:

- $P \xrightarrow{\tau} \sigma Q$  or  $P \xrightarrow{k(\tau)} \sigma Q$ , and  $\langle S; \Sigma \rangle \xrightarrow{*} \langle S'; \Sigma' \rangle$ ;
- $P \xrightarrow{\beta} \epsilon Q$  and  $\langle S; \Sigma \rangle \xrightarrow{\beta}^* \langle S'; \Sigma' \rangle$ .

*Proof.* See Appendix A. □

**Theorem 3.** For all  $P \in \text{Proc}$ ,  $P$  is star simulated by  $\eta(P)$  up to  $\cong_t$ .

## 5 Conclusions and future work

In this paper we have introduced *OCTM*, a higher-order language extending the concurrency model of STM Haskell with composable *open (multi-thread)* transactions. In this language, processes can *join* transactions and transactions can *merge* at runtime. These interactions are driven only by access to shared transactional memory, and hence are implicit and loosely coupled. To this end, we have separated the isolation aspect from atomicity: the `atomic` construct ensures “all-or-nothing” execution but not isolation, while the new constructor `isolated` can be used to guarantee isolation when needed. In order to show the expressive power of *OCTM*, we have provided an adequate implementation in it of *TCCS<sup>m</sup>*, a recently introduced model of open transactions with CCS-like communication. As a side result, we have given a simple typing system for capturing *TCCS<sup>m</sup>* well-formed terms.

Several directions for future work stem from the present paper. First, we plan to implement *OCTM* along the line of STM Haskell, but clearly the basic ideas of *OCTM* are quite general and can be applied to other STM implementations, like C/C++ LibCMT and Java Multiverse.

An interesting possibility is to use *TCCS<sup>m</sup>* as an *exogenous orchestration language* for *OCTM*: the *behaviour* of a transactional distributed system can be described as a *TCCS<sup>m</sup>* term, which can be translated into a *skeleton* in *OCTM* using the encoding provided in this paper; then, the programmer has only to “fill in the gaps”. Thus, *TCCS<sup>m</sup>* can be seen as a kind of “global behavioural type” for *OCTM*.

In fact, defining a proper behavioural typing system for transactional languages like *OCTM* is another interesting future work. Some preliminary experiments have shown that *TCCS<sup>m</sup>* is not enough expressive for modelling the dynamic creation of resources (locations, threads, etc.). We think that a good candidate could be a variant of *TCCS<sup>m</sup>* with local names and scope extrusions, i.e., a “transactional  $\pi$ -calculus”.

Being based on CCS, communication in *TCCS<sup>m</sup>* is synchronous; however, nowadays asynchronous models play an important rôle (see e.g. actors, event-driven programming, etc.). It may be interesting to generalize the discussion so as to consider also this case, e.g. by defining an actor-based calculus with open transactions. Such a calculus can be quite useful also for modelling speculative reasoning for cooperating systems [14–16]. A local version of actor-based open transactions can be implemented in *OCTM* using lock-free data structures (e.g., message queues) in shared transactional memory.

## References

- [1] M. Abadi, A. Birrell, T. Harris, and M. Isard. Semantics of transactional memory and automatic mutual exclusion. *ACM Trans. Program. Lang. Syst.*, 33(1):2, 2011.
- [2] R. Bruni, H. C. Melgratti, and U. Montanari. Nested commits for mobile calculi: Extending join. In J. Lévy, E. W. Mayr, and J. C. Mitchell, editors, *Proc. TCS*, volume 155 of *IFIP*, pages 563–576. Kluwer/Springer, 2004.
- [3] V. Danos and J. Krivine. Transactions in RCCS. In M. Abadi and L. de Alfaro, editors, *Proc. CONCUR*, volume 3653 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 2005.
- [4] E. de Vries, V. Koutavas, and M. Hennessy. Communicating transactions (extended abstract). In P. Gastin and F. Laroussinie, editors, *Proc. CONCUR*, volume 6269 of *Lecture Notes in Computer Science*, pages 569–583. Springer, 2010.

- [5] E. de Vries, V. Koutavas, and M. Hennessy. Liveness of communicating transactions (extended abstract). In K. Ueda, editor, *Proc. APLAS*, volume 6461 of *Lecture Notes in Computer Science*, pages 392–407. Springer, 2010.
- [6] J. Gray and L. Lamport. Consensus on transaction commit. *ACM Transactions Database Systems*, 31(1):133–160, 2006.
- [7] T. Harris, S. Marlow, S. L. P. Jones, and M. Herlihy. Composable memory transactions. In *Proc. PPOPP*, pages 48–60, 2005.
- [8] M. Herlihy and J. E. B. Moss. Transactional memory: Architectural support for lock-free data structures. In A. J. Smith, editor, *Proceedings of the 20th Annual International Symposium on Computer Architecture*, pages 289–300. ACM, 1993.
- [9] M. Herlihy and N. Shavit. Transactional memory: beyond the first two decades. *SIGACT News*, 43(4):101–103, 2012.
- [10] S. L. P. Jones, A. D. Gordon, and S. Finne. Concurrent haskell. In H. Boehm and G. L. S. Jr., editors, *Proc. POPL*, pages 295–308. ACM Press, 1996.
- [11] V. Koutavas, C. Spaccasassi, and M. Hennessy. Bisimulations for communicating transactions - (extended abstract). In A. Muscholl, editor, *Proc. FOSSACS*, volume 8412 of *Lecture Notes in Computer Science*, pages 320–334. Springer, 2014.
- [12] X. Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [13] M. Lesani and J. Palsberg. Communicating memory transactions. In C. Cascaval and P. Yew, editors, *Proc. PPOPP*, pages 157–168. ACM, 2011.
- [14] J. Ma, K. Broda, R. Goebel, H. Hosobe, A. Russo, and K. Satoh. Speculative abductive reasoning for hierarchical agent systems. In J. Dix, J. Leite, G. Governatori, and W. Jamroga, editors, *Proc. CLMAS*, volume 6245 of *Lecture Notes in Computer Science*, pages 49–64. Springer, 2010.
- [15] A. Mansutti, M. Miculan, and M. Peressotti. Multi-agent systems design and prototyping with bigraphical reactive systems. In K. Magoutis and P. Pietzuch, editors, *Proc. DAIS*, volume 8460 of *Lecture Notes in Computer Science*, pages 201–208. Springer, 2014.
- [16] A. Mansutti, M. Miculan, and M. Peressotti. Distributed execution of bigraphical reactive systems. *CoRR*, abs/1503.02434, 2015.
- [17] R. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, 1980.
- [18] N. Shavit and D. Touitou. Software transactional memory. *Distributed Computing*, 10(2):99–116, 1997.

## A Omitted proofs

*Proof of Lemma 1.* The proof proceeds by induction on the syntax of  $TCCS^m$ . We only show three cases:

1. a transition  $P \xrightarrow{\tau_\varepsilon} Q$  resulting from a synchronization outside a transaction;
2. a transition  $P \xrightarrow{k(\tau)} \sigma Q$  resulting from a synchronization inside a transaction;
3. a commit transition  $P \xrightarrow{\text{co}k} Q$ .

1. If  $P \xrightarrow{\tau_\varepsilon} Q$  with (SYNC) rule,  $P = P_1 \mid P_2$ ,  $Q = Q_1 \mid Q_2$ ,  $P_1 \xrightarrow{a} Q_1$ ,  $P_2 \xrightarrow{\bar{a}} Q_2$ .  
 $P_1 = ((\sum_i^{m_1} a_i.R'_i) \mid P'_1) \setminus L_1$  such that  $\exists i. a_i = a$  and  $a \notin L$   
 $P_2 = ((\sum_j^{m_2} b_j.R''_j) \mid P'_2) \setminus L_2$  such that  $\exists j. b_j = \bar{a}$  and  $a \notin L$

$$\begin{aligned} \eta(P) = & \langle ([\text{return } t'_1]_{t_1} \parallel ([\text{return } t_{1m_1}]_{t_{sum1}} \parallel S'_1 \parallel \\ & \parallel ([\xi(a_1, l_{sum1}, R'_1)]_{t_{11}} \parallel \dots \parallel ([\xi(a_{m_1}, l_{sum1}, R'_{m_1})]_{t_{1m_1}} \parallel \\ & \parallel ([\text{return } t'_2]_{t_2} \parallel ([\text{return } t_{2m_2}]_{t_{sum2}} \parallel S'_2 \parallel \\ & \parallel ([\xi(b_1, l_{sum2}, R''_1)]_{t_{21}} \parallel \dots \parallel ([\xi(b_{m_2}, l_{sum2}, R''_{m_2})]_{t_{2m_2}}; (\Theta, \emptyset)) \rangle \end{aligned}$$

From hypothesis, there exists a thread  $t_{1r} \in \{t_{11}, \dots, t_{1m_1}\}$  such that

$$([\varrho(a_r.R'_r)]_{t_{1r}} = ([\text{recv } a \ l_{sum1} \ R'_r]_{t_{1r}})$$

and exists another thread  $t_{2s} \in \{t_{21}, \dots, t_{2m_2}\}$  s.t.

$$([\varrho(a_s.R''_s)]_{t_{2s}} = ([\text{send } a \ l_{sum2} \ R''_s]_{t_{2s}})$$

$l_{sum1}$  and  $l_{sum2}$  are locations created by threads  $t_{sum1}$  and  $t_{sum2}$  from the code generated by encoding of the sums.

$$\begin{aligned} & \langle \dots \parallel ([\text{recv } a \ l_{sum1} \ \varrho(R'_r)]_{t_{1r}} \parallel ([\text{send } a \ l_{sum2} \ \varrho(R''_s)]_{t_{2s}} \parallel \dots; (\Theta, \emptyset)) \rangle \\ \xrightarrow{*} & \langle ([\text{return } t'_1]_{t_1} \parallel ([\text{return } t_{1m_1}]_{t_{sum1}} \parallel S'_1 \parallel \\ & \parallel ([\text{return}]_{t_{11}} \parallel \dots \parallel ([\varrho(R'_r)]_{t_{1r}} \parallel \dots \parallel ([\text{return}]_{t_{1m_1}} \parallel \\ & \parallel ([\text{return } t'_2]_{t_2} \parallel ([\text{return } t_{2m_2}]_{t_{sum2}} \parallel S'_2 \parallel \\ & \parallel ([\text{return}]_{t_{21}} \parallel \dots \parallel ([\varrho(R''_s)]_{t_{2s}} \parallel \dots \parallel ([\text{return}]_{t_{2m_2}}; (\Theta', \emptyset)) \rangle = \langle S; \Sigma \rangle \end{aligned}$$

where  $\Theta'(l_{sum1}) = \text{False}$ ,  $\Theta'(l_{sum2}) = \text{False}$

`recv a lsum1 ρ(R'r)` and `send a lsum2 ρ(R''s)` can in order execute the isolated blocks, and at the end reduce to continuations  $\varrho(R'_r)$  and  $\varrho(R''_s)$ .

Other threads forked by threads  $t_{sum1}, t_{sum2}$  can only reduce to `return` because  $\Theta'(l_{sum1}) = \text{False}$  and  $\Theta'(l_{sum2}) = \text{False}$ : threads  $t_{1r}, t_{2s}$  modified  $l$ -variables through the synchronization code inside isolated blocks. We can observe  $\langle S; \Sigma \rangle \cong_t \eta(Q)$ , in fact  $Q = (R'_i \mid P'_1) \setminus L_1 \mid (R''_j \mid P'_2) \setminus L_2$ ,  $\eta(Q) = \langle S_q; \Sigma_q \rangle$ .

$$\begin{aligned} \eta(Q) = & \langle ([\text{return } t'_1]_{t_1} \parallel ([\varrho(R'_i)]_{t_i} \parallel S'_1 \parallel \\ & \parallel ([\text{return } t'_2]_{t_2} \parallel ([\varrho(R'_j)]_{t_j} \parallel S'_2; (\Theta'_q, \emptyset)) \rangle = \langle S_q; \Sigma_q \rangle \end{aligned}$$

where  $\forall ch \in L_1 \uplus L_2. \Theta'_q(ch) = \text{M0}$

$\langle S_q; \Sigma_q \rangle$  and  $\langle S; \Sigma \rangle$  are different only in local variables and for reduced threads, thus  $\langle S; \Sigma \rangle \cong_t \eta(Q)$ .

2. If  $P \xrightarrow{i,j \mapsto k} Q$  with (TSYNC) rule,  $P = (P_1 \mid P_2)$ ,  $Q = (Q_1 \mid Q_2)$ ,  $P_1 \xrightarrow{i \mapsto k} Q_1$ ,  $P_2 \xrightarrow{j \mapsto k} Q_2$ .

$P = \langle\langle P_1 \triangleright_i C_1 \rangle\rangle \mid \langle\langle P_2 \triangleright_j C_2 \rangle\rangle$  and  $P_1 = P_{11} \mid \dots \mid P_{1m_1}$ ,  $P_2 = P_{21} \mid \dots \mid P_{2m_2}$ .

$$\begin{aligned} \eta(P) = & \langle (\langle P'_1 \triangleright \varrho(C_1) \rangle)_{t_1,i} \parallel (\langle P'_2 \triangleright \varrho(C_2) \rangle)_{t_2,i} \parallel (\langle \varrho(P'_{11}) \triangleright \mathbf{return} \rangle)_{t_{11},i} \parallel \dots \\ & \dots \parallel (\langle \mathbf{recv} \; a \; l \; P'_{1r} \triangleright \mathbf{return} \rangle)_{t_r,i} \parallel \dots \parallel (\langle \varrho(P'_{1m_1}) \triangleright \mathbf{return} \rangle)_{t_{1m_1},i} \parallel \dots \\ & \dots \parallel (\langle \varrho(P'_{21}) \triangleright \mathbf{return} \rangle)_{t_{21},j} \parallel \dots \parallel (\langle \mathbf{send} \; a \; l \; P'_{2s} \triangleright \mathbf{return} \rangle)_{t_s,j} \dots \\ & \dots \parallel (\langle \varrho(P'_{2m_2}) \triangleright \mathbf{return} \rangle)_{t_{2m_2},j}; (\Theta, \Delta) \rangle \\ \rightarrow^* & \langle (\langle P'_1 \triangleright \varrho(C_1) \rangle)_{t_1,j} \parallel (\langle P'_2 \triangleright \varrho(C_2) \rangle)_{t_2,j} \parallel \dots \\ & \dots \parallel (\langle P'_{1r} \triangleright \mathbf{return} \rangle)_{t_r,j} \parallel \dots \parallel (\langle P'_{2s} \triangleright \mathbf{return} \rangle)_{t_s,j}; (\Theta, \Delta') \rangle = \langle S; \Sigma \rangle \end{aligned}$$

$$\begin{aligned} \langle S; \Sigma \rangle \cong_t \eta(Q) = & \\ = & \eta(\langle\langle P_{11} \mid \dots \mid P'_{1r} \mid \dots \mid P_{1m_1} \triangleright_k C_1 \rangle\rangle \mid \\ & \mid \langle\langle P_{21} \mid \dots \mid P'_{2s} \mid \dots \mid P_{2m_2} \triangleright_k C_2 \rangle\rangle) \\ = & \langle (\dots \triangleright \varrho(C_1))_{t_1,k} \parallel (\dots \triangleright \varrho(C_2))_{t_2,k} \parallel \dots \\ & \dots \parallel (\langle P'_{1r} \triangleright \mathbf{return} \rangle)_{t_r,k} \parallel \dots \parallel (\langle P'_{2s} \triangleright \mathbf{return} \rangle)_{t_s,k}; (\Theta_q, \Delta_q) \rangle \end{aligned}$$

3. If  $P \xrightarrow{\text{co}^k} Q$  with (TCO) rule,  $P = \langle\langle R \triangleright_k N \rangle\rangle$  and  $\langle\langle R \triangleright_k N \rangle\rangle \xrightarrow{\text{co}^k} R'$  where  $R' = \prod_{i=1}^m R_i$

$$\begin{aligned} \eta(P) = & \langle (\langle \mathbf{recv} \; \text{co} \; l_t \; \mathbf{return} \triangleright \varrho(N); R \rangle)_{t,k} \parallel (\langle \mathbf{return} \; t_m \triangleright \mathbf{return} \rangle)_{t',k} \parallel \\ & \parallel (\langle \varrho(R_1) \triangleright \mathbf{return} \rangle)_{t_1,k} \parallel \dots \parallel (\langle \varrho(R_j) \triangleright \mathbf{return} \rangle)_{t_j,k} \parallel \dots \\ & \dots \parallel (\langle \varrho(R_m) \triangleright \mathbf{return} \rangle)_{t_m,k}; (\Theta, \Delta) \rangle \end{aligned}$$

where

$$\Delta(l_t) = (\mathbf{True}, k), \quad \Delta(\text{co}) = (\mathbf{M0}, k)$$

$$R = \mathbf{bang} \; \mathbf{psi}$$

$$\exists j \in \{1, \dots, m\}, \quad R_j = \text{co}.R'_j, \quad \varrho(R_j) = \mathbf{send} \; \text{co} \; l \; \varrho(R'_j).$$

$$\begin{aligned} \eta(P) \rightarrow^* & \langle (\langle \mathbf{recv} \; \text{co} \; l_t \; \mathbf{return} \triangleright \varrho(N); R \rangle)_{t,k} \parallel (\langle \mathbf{return} \; t_m \triangleright \mathbf{return} \rangle)_{t',k} \parallel \\ & \parallel \dots \parallel (\langle \mathbf{send} \; \text{co} \; l_{t_j} \; \varrho(R'_j) \triangleright \mathbf{return} \rangle)_{t_j,k} \parallel \dots; (\Theta, \Delta'') \rangle \end{aligned}$$

At this point threads  $t, t_j$  can synchronize through  $\text{co}$  variable and transaction  $k$  can com-

mit.

$$\begin{aligned}
& \rightarrow^* \langle (\text{recv co } l_t \text{ return} \triangleright \varrho(N); \text{bang psi})_{t,k} \parallel \\
& \quad \| (\text{return } t_m \triangleright \text{return})_{t',k} \| \| (\varrho(R_1) \triangleright \text{return})_{t_1,k} \| \dots \\
& \quad \dots \| (\text{send co } l_{t_j} \varrho(R'_j) \triangleright \text{return})_{t_j,k} \| \dots \| (\varrho(R_m) \triangleright \text{return})_{t_m,k}; (\Theta, \Delta'') \rangle \\
& \rightarrow^* \langle (\text{return} \triangleright \varrho(N); \text{bang psi})_{t,k} \parallel \\
& \quad \| (\text{return } t_m \triangleright \text{return})_{t',k} \| \\
& \quad \| (\varrho(R_1) \triangleright \text{return})_{t_1,k} \| \dots \| (\varrho(R'_j) \triangleright \text{return})_{t_j,k} \| \dots \\
& \quad \dots \| (\varrho(R_m) \triangleright \text{return})_{t_m,k}; (\Theta, \Delta''') \rangle \\
& \xrightarrow{\text{co}_k} \langle (\text{bang psi})_t \parallel (\text{return } t_m)_{t,t'} \parallel \\
& \quad \| (\varrho(R_1))_{t_1} \| \dots \| (\varrho(R'_j))_{t_j} \| \dots \| (\varrho(R_m))_{t_m}; \Sigma \rangle = \langle S; \Sigma \rangle \\
& \quad \text{where } \Sigma = \text{commit}(k, \Theta, \Delta''')
\end{aligned}$$

$$\langle S; \Sigma \rangle \cong_t \eta(Q), Q = R' = \prod_1^m R_i \text{ and } \exists j \in \{1, \dots, m\} : R_j = R'_j$$

$$\eta(Q) = \eta(R') = \langle (\varrho(R_1))_{t_1} \| \dots \| (\varrho(R'_j))_{t_j} \| \dots \| (\varrho(R_m))_{t_m}; \Sigma_q \rangle \cong_t \langle S; \Sigma \rangle$$

□

*Proof of Lemma 2.* The proof goes through induction on the semantic of  $TCCS^m$ . Here we show only 3 cases, first when  $P$  are two processes that can perform a synchronization outside transactions, second  $P$  synchronizes inside transactional processes, third a transactional process commits.

1. If  $P = a.P_1 \mid \bar{a}.P_2$  From  $\eta(P)$  we can move to another state of the  $OCTM$  machine  $\langle S; \Sigma \rangle$  and  $\langle S; \Sigma \rangle \cong_t \eta(Q)$ :

$$\begin{aligned}
\eta(P) = & \langle (\text{return } t'_1)_{t_1} \parallel (\text{recv } a \ r_l \ \varrho(P_1))_{t'_1} \parallel \\
& \| (\text{return } t'_2)_{t_2} \parallel (\text{recv } a \ r_l \ \varrho(P_2))_{t'_2}; (\Theta', \emptyset) \rangle \\
\rightarrow^* & \langle (\text{return } t'_1)_{t_1} \parallel (\varrho(P_1))_{t'_1} \parallel (\text{return } t'_2)_{t_2} \parallel (\varrho(P_2))_{t'_2}; (\Theta', \emptyset) \rangle = \langle S; \Sigma \rangle \\
& \text{where } \Theta'(l_{t_1}) = \text{False}, \Theta'(l_{t_2}) = \text{False}, \Theta'(a) = \text{M0}
\end{aligned}$$

If  $P \xrightarrow{\tau} Q$ , then  $Q = P_1 \mid P_2$ ,  $\eta(Q) = \langle (\varrho(P_1))_{t_1} \| (\varrho(P_2))_{t_2}; (\Theta_q, \emptyset) \rangle$ , we can observe  $\langle S'; \Sigma' \rangle \cong_t \eta(Q)$ .

2. If  $P = \langle\langle a.P_1 \triangleright_i Q_1 \rangle\rangle \mid \langle\langle \bar{a}.P_2 \triangleright_j Q_2 \rangle\rangle$ ,  $\eta(P) \rightarrow \langle S; \Sigma \rangle \rightarrow^* \langle S'; \Sigma' \rangle$  the computations are exactly the same as the previous point, but all variables are tentative in  $\Delta$ . It is easy to see that  $P \xrightarrow{\tau(k)}_{i,j \mapsto k} Q$  and  $\langle S; \Sigma \rangle \cong_t \eta(Q)$ .

3. If  $P = \langle\langle co.P' \triangleright_k Q \rangle\rangle$

$$\begin{aligned}
\eta(P) &= \\
&= \langle \langle \text{recv } co \ l \ \text{return} \triangleright \varrho(Q); \text{bang } \text{psi} \rangle \rangle_{t,k} \parallel \\
&\quad \parallel \langle \langle \text{send } co \ l \ \varrho(P') \triangleright \text{return} \rangle \rangle_{t',k}; (\Theta, \Delta) \rangle \\
&\rightarrow \langle \langle M \triangleright \varrho(Q); \text{bang } \text{psi} \rangle \rangle_{t,k} \parallel \\
&\quad \parallel \langle \langle \text{send } co \ l \ \varrho(P') \triangleright \text{return} \rangle \rangle_{t',k}; (\Theta, \Delta') \rangle = \langle S; \Sigma \rangle \\
&\quad \text{where } \Delta'(np) = \text{nonce}_{t_1} \\
&\xrightarrow{* \text{ co}_k} \langle \langle \text{bang } \text{psi} \rangle \rangle_t \parallel \langle \langle \varrho(P') \rangle \rangle_{t'}; (\Theta', \emptyset) \rangle \\
&\quad \text{where } \Sigma' = \text{commit}(k, \Theta', \Delta') \\
&= \langle S'; \Sigma' \rangle
\end{aligned}$$

$$P = \langle\langle co.P' \triangleright_k N \rangle\rangle \xrightarrow{\text{co}_k} P' = Q, \eta(Q) = \langle \langle \varrho(P') \rangle \rangle_{t_q}; \Sigma_q \rangle, \eta(Q) \cong_t \langle S'; \Sigma' \rangle$$

□